# Comparison of 250 MHz R10K Origin 2000 and 400 MHz Origin 2000 Using NAS Parallel Benchmarks

## NAS Technical Report 01-007

**Raymond D. Turney**

**Computer Sciences Corporation**
**Numerical Aerospace Simulation Facility**
**NASA Ames Research Center**
**MS 258-6**
**Moffett Field, CA. 94035**
**email: turney@nas.nasa.gov**

## Introduction

This report describes results of benchmark tests on Steger, a 250 MHz Origin 2000 system with R10K processors, currently installed at the NASA Ames National Advanced Supercomputing (NAS) facility. For comparison purposes, the tests were also run on Lomax, a 400 MHz Origin 2000 with R12K processors.

The BT, LU, and SP application benchmarks in the NAS Parallel Benchmark Suite and the kernel benchmark FT were chosen to measure system performance. Having been written to measure performance on Computational Fluid Dynamics applications, these benchmarks are assumed appropriate to represent the NAS workload. Since the NAS runs both message passing (MPI) and shared-memory, compiler directive type codes, both MPI and OpenMP versions of the benchmarks were used. The MPI versions used were the latest official release of the NAS Parallel Benchmarks, version 2.3. The OpenMP versions used were PBN3b2, a beta version that is in the process of being released. NPB 2.3 and PBN 3b2 are technically different benchmarks, and NPB results are not directly comparable to PBN results.

Links to descriptions of the benchmarks themselves:

NPB description

PBN description

All runs were Class C, and compiled with 64-bit addressing. The MPI programs were compiled with the -O2 compiler flag, described as extensive optimization by the SGI compiler man pages. The OpenMP runs were compiled with the -O3 compiler flag, the highest level of optimization on the Origin. Different flags were used because the MPI BT Class C benchmark ran faster when compiled with the -O2 flag. After running the MPI benchmarks it was discovered that this was not the case for all benchmarks, so the -O3 flag, normally the faster option, was used for the Open MP benchmarks. When the MPI benchmarks were run compiled with the -O3 flag on Steger, the timings were within 5% of the times obtained compiling with -O2, so compiler flag choice did not significantly affect the results.

## Summary of Results

The average execution times of the MPI runs for each machine, as well as the sum of the averages, are listed in Table 1.

Table 1 - Summary of MPI Results

| Machine | BT(sec) | FT(sec) | LU(sec) | SP(sec) | Total |
|---|---|---|---|---|---|
| Steger - R10K O2K 250MHz | 2925.58 | 631.09 | 1255.93 | 1894.41 | 6718.01 |
| Lomax - R12K O2K 400MHz | 2613.65 | 430.73 | 922.04 | 1433.78 | 5400.20 |

The ratio of the total times is:

6718.01/5400.20 = 1.24

Averaging over all the MPI benchmark runs, Lomax was about a quarter faster than Steger.

Table 2 lists the average execution times of the OpenMP runs for each machine, and the sum of the averages.

Table 2 - Summary of OpenMP Results

| Machine | BT(sec) | FT(sec) | LU(sec) | SP(sec) | Total(sec) |
|---|---|---|---|---|---|
| Steger R10K 250MHz | 1354.28 | 276.88 | 1503.03 | 1431.55 | 4565.74 |
| Lomax - O2K 400MHz | 969.44 | 297.41 | 996.88 | 1178.87 | 3442.60 |

The ratio of the total times is:

4565.74/3442.60 = 1.33

Averaging over all the benchmarks, Lomax was about a third faster on the OpenMP version of the benchmarks.

These results, MPI and OpenMP, suggest that codes represented by NPB and PBN Class C, should run somewhere between 25 and 33 percent faster on Lomax. No explanation for the oddity that OpenMP FT actually ran faster on steger than on lomax is currently available.

## Results

Class C was used because it is the largest size of these benchmarks, and the primary purpose of these machines is to run large jobs. Sixteen cpus was selected as representative of a small to medium size job, and a convenient number for the benchmarks and the machines. The O2 optimization level was selected for the MPI benchmarks because the MPI NPB BT Class C runs faster on the O2K when compiled with O2 optimization than with O3 optimization. The O3 optimization level was selected for the OpenMP benchmarks because it is the highest level of optimization on the Origin. 64-bit addressing was selected because it is impossible to compile Class C for some of these benchmarks using 32-bit addresses.

The mean, median, and standard deviation of the timings and MOPS counts for each benchmark are presented below in Tables 3, 4, 5 and 6. To get a reasonable sample, seven runs of each benchmark were done on each machine.

All runs were done on a machine controlled by a custom PBS scheduler written by Ed Hook, of CSC Corp working at the NASA NAS division, which uses cpusets and an awareness of machine topology to insure memory allocation and execution on physically contiguous nodes. Because the machines were space shared, not time shared, interference from other jobs was minimized. The lack of run time variation among benchmark runs supports this hypothesis.

### Table 3 - Steger - 250 MHz R10K - MPI results

| **BT Class C** | **Seconds** | **MOPS** |
|---|---|---|
| Median | 2922.24 | 980.85 |
| Mean | 2925.58 | 979.75 |
| Std. Dev. | 120.77 | 4.77 |
| **FT Class C** | **Seconds** | **MOPS** |
| Median | 623.53 | 635.72 |
| Mean | 631.09 | 628.81 |
| Std. Dev. | 22.90 | 20.89 |
| **LU Class C** | **Seconds** | **MOPS** |

| | | |
|---|---|---|
| Median | 1266.25 | 1610.26 |
| Mean | 1266.93 | 1609.40 |
| Std. Dev. | 2.38 | 3.02 |
| **SP Class C** | **Seconds** | **MOPS** |
| Median | 1894.96 | 765.24 |
| Mean | 1894.41 | 765.47 |
| Std. Dev. | 4.23 | 1.71 |

Hardware info:
IRIX64 steger 6.5 10120851 IP27
256 250 MHZ IP27 Processors
CPU: MIPS R10000 Processor Chip Revision: 3.4
FPU: MIPS R10010 Floating Point Chip Revision: 3.4
Main memory size: 65536 Mbytes
Instruction cache size: 32 Kbytes
Data cache size: 32 Kbytes
Secondary unified instruction/data cache size: 4 Mbytes

## Table 4 - Lomax - 400 MHz R12K Origin 2000 MPI Results

| **BT Class C** | **Seconds** | **MOPS** |
|---|---|---|
| Median | 2602.80 | 1101.23 |
| Mean | 2613.65 | 1096.90 |
| Std. Dev. | 26.31 | 10.93 |
| **FT Class C** | **Seconds** | **MOPS** |
| Median | 427.17 | 927.95 |
| Mean | 430.73 | 921.48 |
| Std. Dev. | 17.24 | 35.71 |
| **LU Class C** | **Seconds** | **MOPS** |
| Median | 912.10 | 2235.50 |
| Mean | 922.04 | 2212.27 |
| Std. Dev. | 20.17 | 46.77 |
| **SP Class C** | **Seconds** | **MOPS** |
| Median | 1409.46 | 1028.83 |

| | | |
|---|---|---|
| Mean | 1433.78 | 1012.24 |
| Std. Dev. | 46.36 | 31.42 |

Hardware info:
IRIX64 lomax 6.5 07061118 IP27
Processors: 512 400 MHZ IP27 Processors
CPU: MIPS R12000 Processor Chip Revision: 3.5
FPU: MIPS R12010 Floating Point Chip Revision: 0.0
Main memory size: 196608 Mbytes
Instruction cache size: 32 Kbytes
Data cache size: 32 Kbytes
Secondary unified instruction/data cache size: 8 Mbytes

## Table 5 - Steger - 250 MHz R10K OpenMP

>
>

| **BT Class C** | **Seconds** | **MOPS** |
|---|---|---|
| Median | 1365.21 | 2099.52 |
| Mean | 1354.28 | 2117.47 |
| Std. Dev. | 31.68 | 50.31 |
| **FT Class C** | **Seconds** | **MOPS** |
| Median | 278.71 | 1421.48 |
| Mean | 276.88 | 1432.07 |
| Std. Dev. | 5.36 | 28.03 |
| **LU Class C** | **Seconds** | **MOPS** |
| Median | 1494.65 | 1364.20 |
| Mean | 1503.03 | 1356.88 |
| Std. Dev. | 23.97 | 21.15 |
| **SP Class C** | **Seconds** | **MOPS** |
| Median | 1422.39 | 1019.48 |
| Mean | 1431.55 | 1013.26 |
| Std. Dev. | 28.11 | 18.81 |

Hardware info:

IRIX64 steger 6.5 10120851 IP27
256 250 MHZ IP27 Processors
CPU: MIPS R10000 Processor Chip Revision: 3.4
FPU: MIPS R10010 Floating Point Chip Revision: 3.4
Main memory size: 65536 Mbytes
Instruction cache size: 32 Kbytes
Data cache size: 32 Kbytes
Secondary unified instruction/data cache size: 4 Mbytes

## Table 6 - Lomax - 400 MHz R12K OpenMP

| **BT Class C** | **Seconds** | **MOPS** |
|---|---|---|
| Median | 967.25 | 2963.32 |
| Mean | 969.44 | 2956.64 |
| Std. Dev. | 2.63 | 9.66 |
| **FT Class C** | **Seconds** | **MOPS** |
| Median | 297.57 | 1332.08 |
| Mean | 297.41 | 1332.85 |
| Std. Dev. | 1.93 | 8.65 |
| **LU Class C** | **Seconds** | **MOPS** |
| Median | 996.52 | 2046.11 |
| Mean | 996.88 | 2045.37 |
| Std. Dev. | 2.03 | 4.11 |
| **SP Class C** | **Seconds** | **MOPS** |
| Median | 1178.71 | 1230.25 |
| Mean | 1178.87 | 1230.08 |
| Std. Dev. | 1.55 | 1.62 |

Hardware info:
IRIX64 lomax 6.5 07061118 IP27
Processors: 512 400 MHZ IP27 Processors
CPU: MIPS R12000 Processor Chip Revision: 3.5
FPU: MIPS R12010 Floating Point Chip Revision: 0.0
Main memory size: 196608 Mbytes
Instruction cache size: 32 Kbytes
Data cache size: 32 Kbytes
Secondary unified instruction/data cache size: 8 Mbytes

## Related Work

Sheila Faulkner, of CSC Corp at the NASA NAS division, has also benchmarked the Origin 2000 and the Origin 3000 using NAS Parallel Benchmarks. She ran all the NPB MPI benchmarks, and investigated scaling. She did not use the OpenMP versions of the benchmarks, and her Origin 2000 numbers are for turing, a 195 Mhz Origin 2000 no longer used as a compute server at the NAS.

## Future Work

This is the second of what is hoped to be a number of benchmark studies comparing various IPG and NAS hosts to each other.